Computing the k Nearest-Neighbors for all Vertices via Dijkstra

Sariel Har-Peled*

July 27, 2016

arXiv:1607.07818v1 [cs.DS] 26 Jul 2016

Abstract

We are given a directed graph $G = (V, E)$ with n vertices and m edges, with positive weights on the edges, and a parameter $k > 0$. We show how to compute, for every vertex $v \in V$, its k nearest-neighbors. The algorithm runs in $O(k(n \log n + m))$ time, and follows by a somewhat careful modification of Dijkstra's shortest path algorithm.

This result is probably folklore, but we were unable to find a reference to it – thus, this note.

1 The problem

Let $G = (V, E)$ be an undirected graph with n vertices, m edges, and with positive weights on the edges. For $u, v \in V$, let $d_G(u, v)$ denote the shortest path distances in G, and assume the sake of the simplicity of exposition that all the non-trivial shortest path distances in the graph are distinct.

A vertex $u = \nu_i(v)$ is the ***ith nearest-neighbor*** to v, if one partition V into three disjoint sets $C, \{u_i\}, F$, such that (i) $|C| = i - 1$, and (ii) $\forall c \in C$, $\forall f \in F$ we have $d_G(c, v) < d_G(u, v) < d_G(f, v)$. Observe that $\nu_0(v) = v$. For any integer k, the ***k nearest-neighbors*** to v, are the members of the set

$$N_{\le k}(v) = \Big\{ \nu_0(v), \nu_1(v), \ldots, \nu_{k-1}(v) \Big\};$$

that is, they are the k vertices in G that are closest to v.

Our purpose here is to compute for every vertex $v \in V$, the set $N_{\le k}(v)$; that is, to compute for v the k distinct vertices closest to it.

2 A simple (but slower) randomized algorithm

2.1 Algorithm

Let $t = O(k \log n)$, and let $R_1, \ldots, R_t$ be random samples from V, where every vertex is picked into the ith random sample with probability $p = 1/k$. Now, compute for every vertex $v \in V$ its distance from its nearest-neighbor in R_i, for all i. For a specific i, this can be done by performing Dijkstra in G starting from all the vertices of R_i simultaneously (i.e., we create a fake source vertex s, add it to G, and add edges of weight 0 from s to all the vertices of R_i, and perform regular Dijkstra from s). For

*Department of Computer Science; University of Illinois; 201 N. Goodwin Avenue; Urbana, IL, 61801, USA; sariel@illinois.edu; http://sarielhp.org/. Work on this paper was partially supported by a NSF AF awards CCF-1421231, and CCF-1217462.

a vertex v, let $D'(v)$ be its t candidate nearest-neighbors computed by these t executions of Dijkstra. Using hashing, remove duplicates in $D'(v)$ (i.e., a vertex s might be the nearest-neighbor for v in several of these executions – note however that in such a case it is always the same distance). Now, compute the k vertices with the smallest numbers associated with them in $D'(v)$, and let $D(v)$ be the resulting set of nearest-neighbors. We claim that, for all v, the sets $D(v)$ are the desired k nearest-neighbors.

2.2 Analysis

For the running time, observe that performing Dijkstra shortest path algorithm t times takes $O(t(n + m \log n))$ time. All the other work is dominated by this.

As for correctness, we need to argue that if u is the jth nearest-neighbor to v, for $j \leq k$, then $u \in D(v)$. To this end, consider the event $\mathcal{E}_i$ that $u \in R_i$, and none of the vertices of $N_{\leq j-1}(v)$ are in R_i. We have that

$$\alpha = \mathbf{Pr}[\mathcal{E}_i] = p(1-p)^{j-1} \geq (1-1/k)^k/k \geq \exp(-1/2k)^k/k \geq 1/10k.$$

If this happens, then $u \in D'(v)$, which in turn implies that $u \in D(v)$, as desired. The probability that none of the events $\mathcal{E}_1, \ldots, \mathcal{E}_t$ happens is $(1-\alpha)^t \leq (1-1/10k)^t < 1/n^c$ by making t sufficiently large, where c is an arbitrary constant. Since there are nk pairs of (u, v) such that u is one of the k nearest neighbors to v, it follows that the probability this algorithm fails is at most $1/n^{c-2}$.

2.3 The result

Lemma 2.1. *Given a directed graph $G = (V, E)$ with n vertices and m edges, with positive weights on the edges, and a parameter $k > 0$, one can compute, in $O((n \log n + m)k \log n)$ time, for every vertex $v \in V$, its k nearest-neighbors in G. The algorithm succeeds with high probability.*

The above randomized algorithm is inspired by the Clarkson-Shor technique [CS89] and this trick is useful in many other scenarios.

3 A faster algorithm

3.1 The algorithm – a first try

We are going to run (conceptually) n copies of the shortest-path algorithm simultaneously. In particular, let A_v be the shortest-path algorithm starting from the vertex v, for all $v \in V$. We use global heap for the events for all these algorithms together. Here, every event of A_v would be indexed by the source vertex v associated with this algorithm. The algorithm is going to maintain for each vertex the set of $D'(v)$ of nearest-neighbors found so far, and a count $c_v = |D'(v)|$.

Now, when the algorithm extract the next vertex to be visited (i.e., the one of the with lowest candidate distance), we get a triple (v, s, d), where v is the vertex to be handled, s is the source vertex, and d is the proposed distance. Using a hash table, we check in constant time whether v has s as one of its computed nearest neighbors (i.e., check if $s \in D'(v)$), and if so the algorithm continues to the next iteration. Otherwise, the algorithm

(i) adds s to $D'(v)$,

(ii) increase c_v, and

(iii) perform the standard relax operation from v for all the outgoing edges of v (these operations are "marked" by the source vertex s they are being done for).
Specifically, consider an edge $v \to z$ being relaxed, during the handling of the event (v, s, d). The new associated event is $(z, s, d + w(v \to z))$, where $w(v \to z)$ is the weight of the edge $v \to z$. If $s \notin D'(z)$ the algorithm "schedule" this event by inserting it to the heap, otherwise it ignores it.

The basic observation is that once a vertex was visited by k of these parallel executions, it is no longer needed, and it can be "disabled" blocking it from being visited by any other Dijkstra (i.e., as soon as $c_v = k$). From this point on, the algorithm ignores update operations for triplets of the form $(v, \cdot, \cdot)$.

3.1.1 Analysis

Lemma 3.1. *The algorithm computes correctly, for each vertex $v \in V$, its k nearest-neighbors in G.*

Proof: By induction on the distance being computed. For $\ell = 0$, for each vertex $v \in V$, the vertex v is its own 0th nearest-neighbor, of distance 0, and it was computed correctly.

Assume all the relevant distances $< \ell$ were computed correctly (they are $0, 1, \ldots, k$ nearest-neighbor distances for some pairs of vertices in the graph).

So consider a vertex v, where its ith nearest-neighbor is s, let its (real) shortest path from s to v be $\pi = s = v_0 v_1 v_2 \ldots v_t = v$, and assume the length of π is ℓ.

The key observation is that for all the vertices on π, the vertex s must be one of their first i nearest-neighbors. Indeed, if not, then there exists i vertices in G that are closer to v than s, which is a contradiction.

Thus, by induction all the relevant distances for $v_0, v_1, \ldots, v_{t-1}$ were computed correctly by the algorithm, and as such, this path would be considered by the algorithm. Namely, the ith distance to v would be set correctly. ∎

Running time. Since for a vertex v, the counter c_v can be increased only k times, it follows that every edge participates in k relax operations. As such, the total number of relax operations handled by this algorithm is $O(km)$, and this bounds the maximum size of the global heap. The global heap might perform $O(km)$ extract-min operations on the global heap. As such, the overall running time of this algorithm is $O(km \log n)$.

3.2 Speeding up the algorithm

Now, our purpose is to improve the running time – this requires some cleverness with the data-structures being used.

Our first task is to avoid having extract-min operations in the global heap involving a vertex v such that $c_v = k$. To this end, for each vertex $v \in V$, the algorithm maintains a separate queue Q_u that handles all the events for u. Every source vertex $s \in V$ might maintain at most one value at Q_u. Every such queue would maintain its current minimum, and would update it in the global queue Q as necessary (all of these operations are either `insert` or `decrease_key`, and both operations can be done in constant time using the standard Fibonacci heap). That is, the global queue contains at most n values, potentially one from each vertex queue. In particular, once $c_v \geq k$, we disable the queue Q_v, and it not longer participates in the global queue, or takes updates to values in its queue.

3.3 Analysis

Lemma 3.2. *The algorithm running time is* $O(k(n \log n + m))$.

Proof: Every vertex distance is going to be set by at most k of these "parallel" executions of Dijkstra. Whenever it happens, the algorithm performs a relax operation on all the adjacent edges. It follows, that an edge would be relaxed $O(k)$ times. The cost of a relax operation is $O(1)$, and thus the overall cost of these is $O(km)$. Setting one of the k shortest distance values of a vertex takes $O(\log n)$ time, as it involves `extract_min` from the global heap. As such, each vertex would require overall $O(k \log n)$ time for its `extract_min` operations from the global heap and the local heap of this vertex. The bound on the running time now follows readily. ■

3.4 The result

Theorem 3.3. *Given a directed graph* $G = (V, E)$ *with* n *vertices and* m *edges, with positive weights on the edges, and a parameter* $k > 0$*, one can compute, in* $O(k(n \log n + m))$ *time, for every vertex* $v \in V$*, its* k *nearest-neighbors in* G.

The algorithm uses hashing (i.e., randomization), a deterministic version of the algorithm runs in time $O(kn \log n + km \log k)$.

Proof: For the deterministic running time, replace the hash table maintained by each vertex v (which stores all the nearest-neighbors to v discovered so far) by a balanced binary tree. The key observation is that such a tree, for a node v, needs to maintain only the k smallest candidate distances offered to this vertex. In particular, if a new candidate distance is larger than all these k values, we can immediately reject it. Similarly, after insertion we reject any value larger than k smallest values in this data-structure.

As such, each binary tree stores at most $O(k)$ elements, so every basic operation takes $O(\log k)$ time. The algorithm performs $O(k(n + m))$ operations on these lookup data structures, which implies the claimed running time. ■

Remark 3.4. Consider the settings of Theorem 3.3, but in addition there is a set of terminals $T \subseteq V$. One can compute for each vertex of G the k closest *terminals* to it in G using the algorithm of Theorem 3.3. The only modification being that we start the Dijkstra process only from the vertices of T. The running time of the modified algorithm is the same.

Acknowledgments. The author thanks Chandra Chekuri and David Eppstein for useful comments on this note. In addition, Vivek Madan came up independently with a similar solution to the one described in this note.

References

[CS89] K. L. Clarkson and P. W. Shor. Applications of random sampling in computational geometry, II. *Discrete Comput. Geom.*, 4:387–421, 1989.